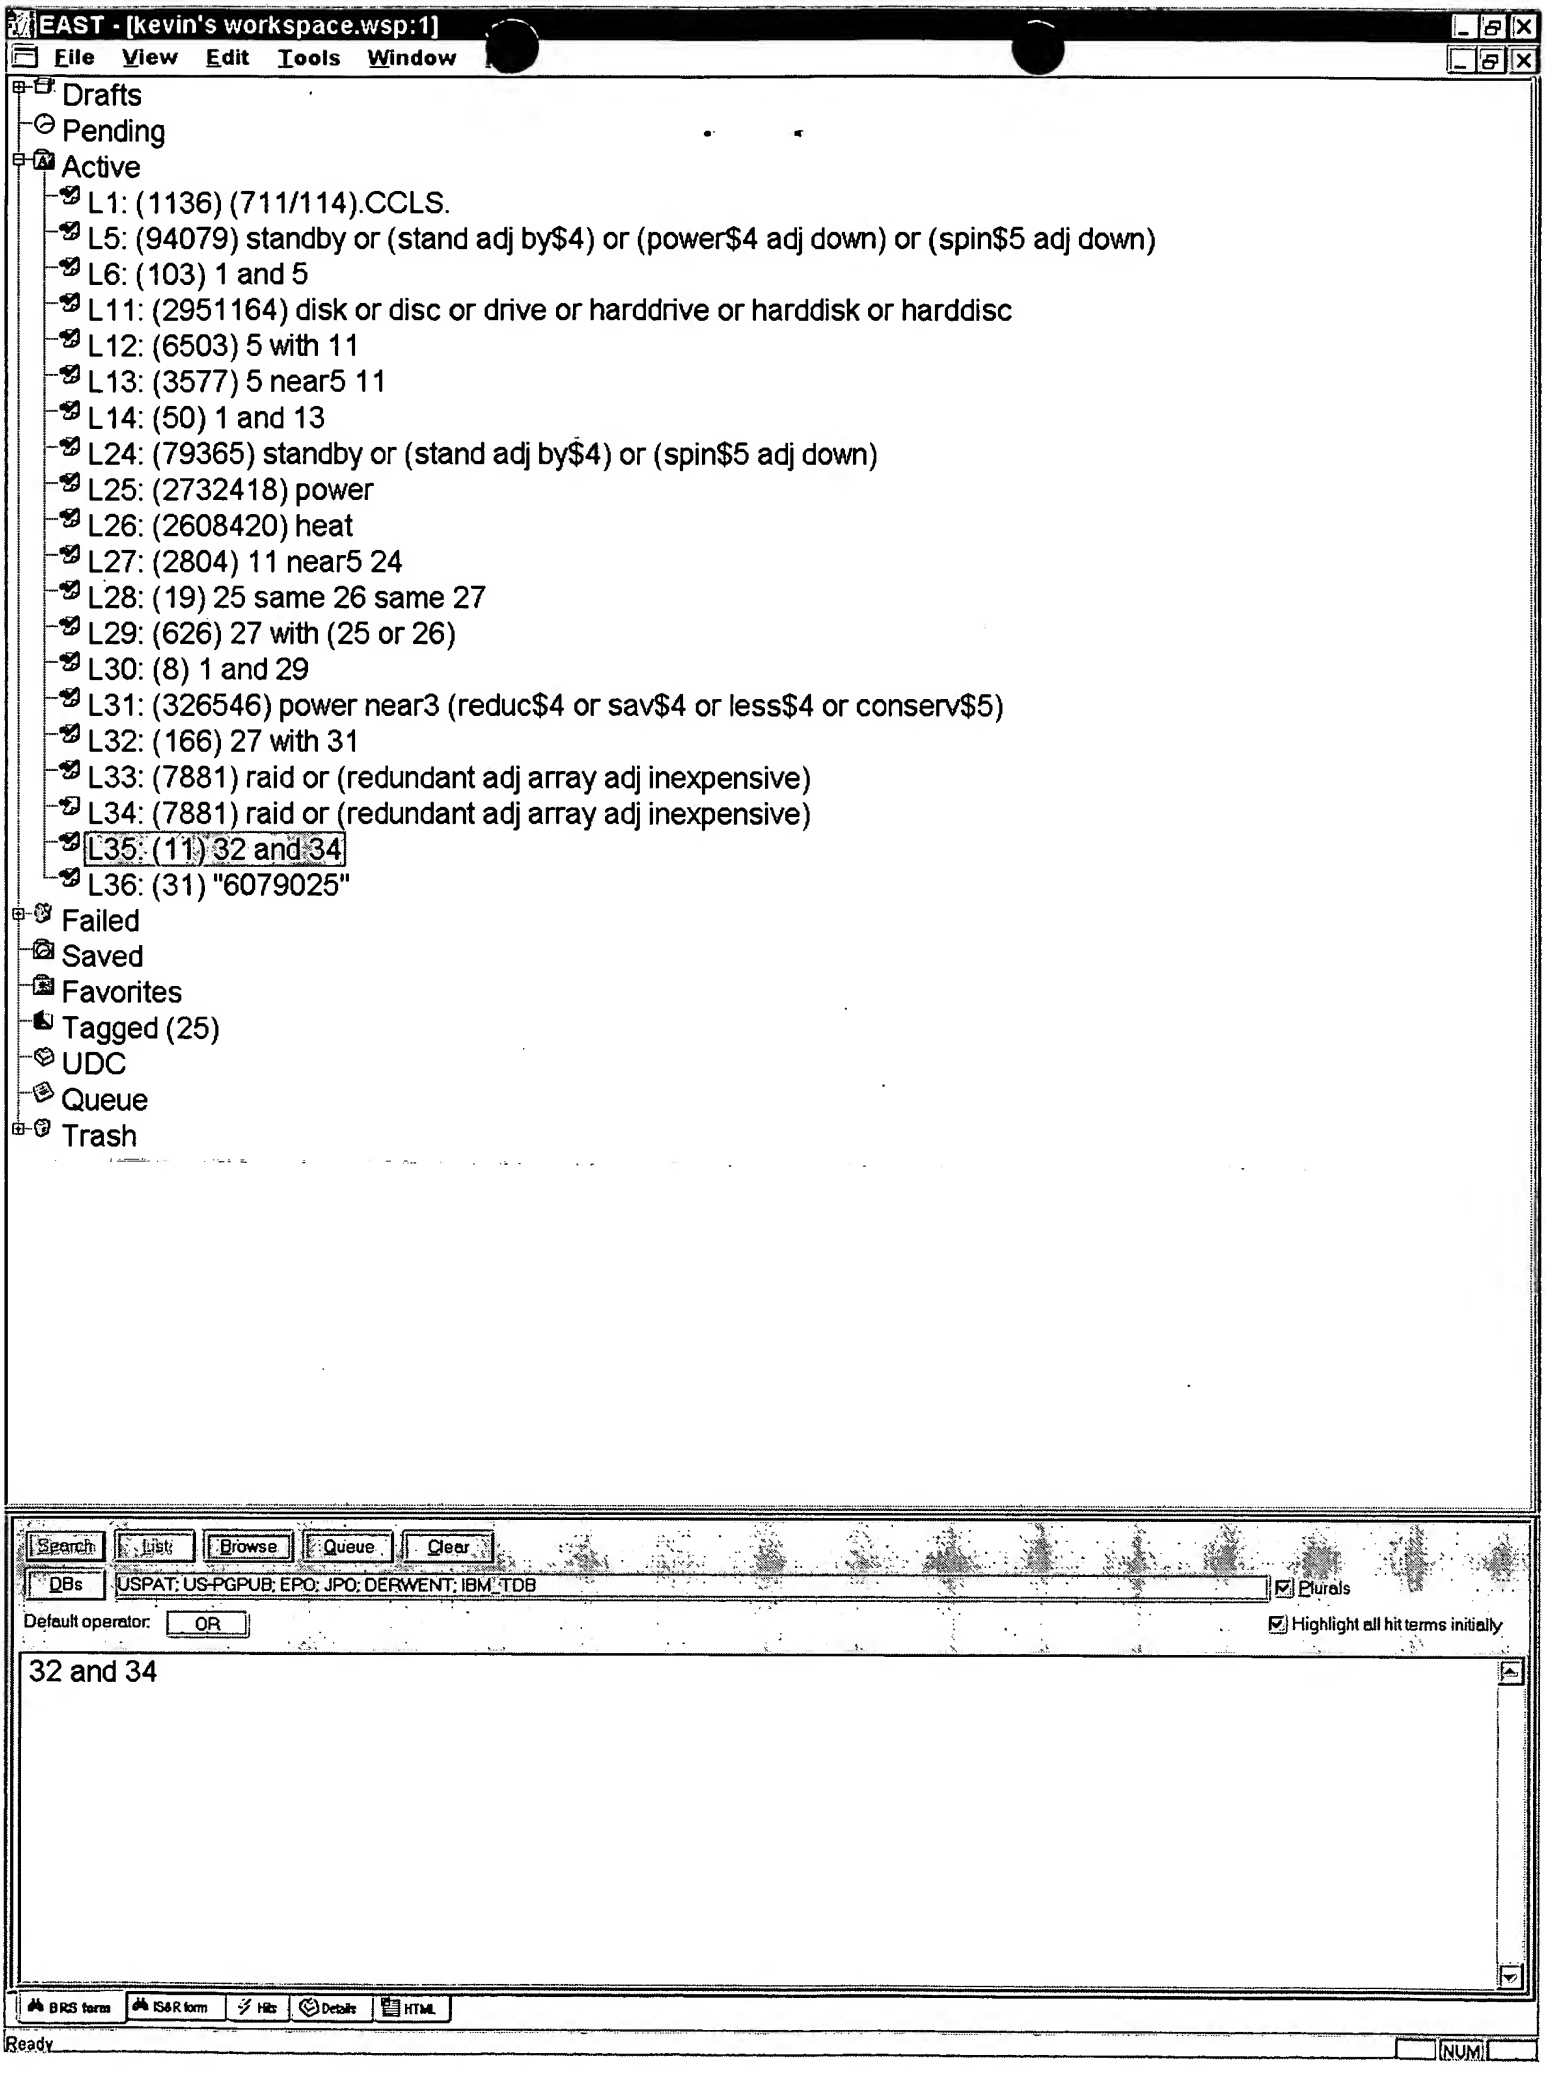
EAST - [kevin's workspace.wsp:1]
File View Edit Tools Window
Drafts
Pending
Active
L1: (1136) (711/114).CCLS.
L5: (94079) standby or (stand adj by$4) or (power$4 adj down) or (spin$5 adj down)
L6: (103) 1 and 5
L11: (2951164) disk or disc or drive or harddrive or harddisk or harddisc
L12: (6503) 5 with 11
L13: (3577) 5 near5 11
L14: (50) 1 and 13
L24: (79365) standby or (stand adj by$4) or (spin$5 adj down)
L25: (2732418) power
L26: (2608420) heat
L27: (2804) 11 near5 24
L28: (19) 25 same 26 same 27
L29: (626) 27 with (25 or 26)
L30: (8) 1 and 29
L31: (326546) power near3 (reduc$4 or sav$4 or less$4 or conserv$5)
L32: (166) 27 with 31
L33: (7881) raid or (redundant adj array adj inexpensive)
L34: (7881) raid or (redundant adj array adj inexpensive)
L35: (11) 32 and 34
L36: (31) "6079025"
Failed
Saved
Favorites
Tagged (25)
UDC
Queue
Trash
Search
List
Browse
Queue
Clear
DBs
USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB
Plurals
Default operator:
OR
Highlight all hit terms initially
32 and 34
BRS form
IS&R form
Hits
Details
HTML
Ready
NUM

L Number	Hits	Search Text	DB	Time stamp
1	1136	(711/114).CCLS.	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/02/10 10:45
5	94079	standby or (stand adj by$4) or (power$4 adj down) or (spin$5 adj down)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/02/10 16:17
6	103	((711/114).CCLS.) and (standby or (stand adj by$4) or (power$4 adj down) or (spin$5 adj down))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/02/10 13:16
11	2951164	disk or disc or drive or harddrive or harddisk or harddisc	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/02/10 15:23
12	6503	(standby or (stand adj by$4) or (power$4 adj down) or (spin$5 adj down)) with (disk or disc or drive or harddrive or harddisk or harddisc)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/02/10 15:24
13	3577	(standby or (stand adj by$4) or (power$4 adj down) or (spin$5 adj down)) near5 (disk or disc or drive or harddrive or harddisk or harddisc)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/02/10 15:25
14	50	((711/114).CCLS.) and ((standby or (stand adj by$4) or (power$4 adj down) or (spin$5 adj down)) near5 (disk or disc or drive or harddrive or harddisk or harddisc))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/02/10 15:25
24	79365	standby or (stand adj by$4) or (spin$5 adj down)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/02/10 16:17
25	2732418	power	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/02/10 16:17
26	2608420	heat	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/02/10 16:17
27	2804	(disk or disc or drive or harddrive or harddisk or harddisc) near5 (standby or (stand adj by$4) or (spin$5 adj down))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/02/10 16:17
28	19	power same heat same ((disk or disc or drive or harddrive or harddisk or harddisc) near5 (standby or (stand adj by$4) or (spin$5 adj down)))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/02/10 16:37
29	626	((disk or disc or drive or harddrive or harddisk or harddisc) near5 (standby or (stand adj by$4) or (spin$5 adj down))) with (power or heat)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/02/10 16:42

30	8	((711/114).CCLS.) and (((disk or disc or drive or harddrive or harddisk or harddisc) near5 (standby or (stand adj by$4) or (spin$5 adj down))) with (power or heat))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/02/10 16:40
31	326546	power near3 (reduc$4 or sav$4 or less$4 or conserv$5)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/02/10 16:42
32	166	((disk or disc or drive or harddrive or harddisk or harddisc) near5 (standby or (stand adj by$4) or (spin$5 adj down))) with (power near3 (reduc$4 or sav$4 or less$4 or conserv$5))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/02/10 16:48
33	7881	raid or (redundant adj array adj inexpensive)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/02/10 16:48
34	7881	raid or (redundant adj array adj inexpensive)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/02/10 16:48
35	11	(((disk or disc or drive or harddrive or harddisk or harddisc) near5 (standby or (stand adj by$4) or (spin$5 adj down))) with (power near3 (reduc$4 or sav$4 or less$4 or conserv$5))) and (raid or (redundant adj array adj inexpensive))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/02/10 17:02
36	31	"6079025"	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/02/10 17:02

DOCUMENT-IDENTIFIER: US 20020144057 A1

TITLE: Archival data storage system and method

-------- KWIC --------

Summary of Invention Paragraph - BSTX (5):

[0004] Magnetic disks are typically used as primary storage for information infrastructures and as storage drives in personal computers, laptop computers, servers, and the like. A number of power saving techniques have been proposed for laptop computers. Software controlled power saving modes have been used to control power consumption during periods of inactivity. Adaptive algorithms which analyze access patterns to adaptively determine when to spin disks up or down to reduce power consumption. Such algorithms, however, usually focus on reducing the power consumption of laptop computers whose disks are specifically designed to spin up and spin down more times than required during the typical life expectancy of a laptop computer. Disks for desktops or servers are usually engineered to handle a limited number of starts and stops. Applying the same power conservation methods used with laptop computers to disk-based archival systems would shorten disk lifetime. Furthermore, these power saving techniques do not address the problem of checking or maintaining the integrity of data stored on disks for extended periods of time.

US-PAT-NO: 6076142

DOCUMENT-IDENTIFIER: US 6076142 A

TITLE: User configurable raid system with multiple data bus segments and removable electrical bridges

-------- KWIC --------

US Patent No. - PN (1):
6076142

Detailed Description Text - DETX (17):
In the RAID mode, when the host sends data to be written to a redundancy group (i.e., a RAID set), the controller 24 stores the data in its cache prior to writing the data to the disk drives. The cache is flushed to make room for additional write data only after the currently stored data has been written onto one or more of the drives in accordance with the RAID level set. In the event of a power failure, the data in cache will be retained under power from the backup battery 29 for a period of time (e.g. 48 hours) to allow adequate time for the system administrator to restore the power. When drive module failure occurs, the system automatically switches over from the defective drive module to a standby "hot spare" drive module that is already installed and operational and not currently configured for use in the system. The system automatically initiates a rebuild operation to reconstruct the failed drive's data onto the hot spare drive module. The rebuild operation by itself is also known in the art and does not form a part of the present invention. Specific structural and functional components and features are described below.

Detailed Description Text - DETX (20):
Once fully inserted in the bays, the modules 14 are powered up (the hard drive may be spinning at this state) and are latched in the bays of the housing 11 by an electromechanical latch under software control. To remove a module, it is first powered down or the hard drive is spun down before the latch can be released, either under software control or by a user actuating a switch 34 (FIG. 5) on the housing 11. This latch control feature has been described in detail in the copending patent application Ser. No. 08/618,290, now abandoned and also the DataDock Storage Systems User's Guide published by the assignee of the present invention, MicroNet Technology, Inc. (June 1996; Part No. MA 1036, Rev. 3) which is fully incorporated by reference herein. A manual release mechanism is provided as a backup to the electronic control. The novel features of this manual release will be described in detail below.

Detailed Description Text - DETX (22):
Referring to FIGS. 6 to 8, each drive module 14 is latched in place in a module bay 111 by a latch 110 against a notch 112 in the module housing 15 when

coupled at hinge 310 to the latch 110. One can push the release button 80, thereby pushing the lever 308, to pivot the latch 110 away from the module housing 15 to a position shown in dotted line 115 in FIG. 7. The latch 110 is disengaged from the module housing 15 to allow the latter to be pulled out of its bay by the user. A spring 320 supported by a stop 322 on the rod 116 biases the rod 116 and its push button 80 back to its original position. It is noted that when the manual release is being used, precaution should be exercised to wait for the drive module 14 to be powered down first.

Detailed Description Text - DETX (48):

The DRIVE LED 50 will illuminate, an audible alarm will sound and the SCSI ID LED 46 will blink when a drive module 14 in a RAID mode has failed. The RAID system 10 will automatically spin down, power off and only unlatch the failed disk module. This prevents further damage to the module and allows the administrator to easily identify which module has failed. This is one of the more important aspects of the present invention. The present system effectively prevents a wrong drive from being removed in the event of a power failure. A more detailed discussion of this will follow under the section "Disk Failure Detection".

Detailed Description Text - DETX (49):

The TEMP Alarm LED 41 will illuminate and an audible alarm will sound when the temperature of the RAID system, an installed drive module, or the room has exceeded normal levels. In the illustrated embodiment, there are a total of nine temperature sensors 54 in the system. One sensor which is located at below the power switch 47 on the front panel (at 39, FIG. 5) monitors the ambient temperature outside the system housing. Another sensor is locate at the backplane 21 for monitoring the temperature within the housing. In each drive module, there is a temperature sensor which monitors any heat buildup within the drive modules. Two threshold levels are preset for alerting to an overheating condition in the modules. When the module temperature reaches the first level, the TEMP Alarm LED will blink and the audible alarm will sound. The problem should be corrected or the system should be shut down. If the temperature climbs to the second level, the RAID system will spin down and turn off all the drives to protect them from damage. The TEMP Alarm LED will then turn on continuously and stop blinking.

Detailed Description Text - DETX (61):

Mode 1 (semiautomatic mode)--Module power down and release are initiated using the drive buttons 34 on the front panel. To remove the module, the operator presses the drive button and the power down sequence is initiated.

Detailed Description Text - DETX (62):

During this time, the button blinks to indicate that the power down sequence is in process. When it is completed, the LED goes off, the latch is released, and the module 14 can be removed.

US-PAT-NO: 6097679

DOCUMENT-IDENTIFIER: US 6097679 A
See image for Certificate of Correction

TITLE: Method for controlling disk-type storage device

--------- KWIC --------

Brief Summary Text - BSTX (7):
The purpose of controlling the disk storage device by application of the idle, standby, and sleep modes is to prevent unnecessary power consumption and heat generation, and to promptly respond to process commands from the data processing unit. However, the conventional method of selecting between the idle mode, that rotates the disk drive motor at normal speed, and the standby mode, that requires the restarting of the disk drive motor when a process command is issued, is inadequate in preventing unnecessary power consumption and heat generation.

Brief Summary Text - BSTX (12):
Therefore, the method of the present invention is characterized by an intermediate standby mode set between the idle and standby modes of the conventional method for controlling a disk-type storage device. The first reference time period for transitioning from the idle mode to the intermediate mode is set shorter than the reference standby interval in the conventional method. The intermediate mode follows the idle mode which rotates the disk at a predetermined revolution speed, and precedes the standby mode which stops the rotation of the disk. Accordingly, since the intermediate standby mode continues to rotate the disk, although at a slower revolution speed than that of the idle mode, the present invention advantageously provides effective suppression of power consumption and heat generation, while allowing prompt responses to new process commands. This is achieved without having to restart the disk driving motor as required by the conventional method's transition directly from the idle mode to the standby mode.

US-PAT-NO: 5459853

DOCUMENT-IDENTIFIER: US 5459853 A

TITLE: Efficient variable-block data storage system employing a staggered fixed-block-architecture array

-------- KWIC --------

Detailed Description Text - DETX (40):

Synchronous DASD arrays often employ parity techniques to recover from single DASD failures without interrupting system operation. Normally such a parity technique employs an additional DASD (N+1) for every N DASDs in an array. A single block on the (N+1).sup.th parity DASD contains the parity of the data in a corresponding sequence of blocks on the remaining N DASDs in the array. When a DASD fails, data from the surviving DASDs and the parity DASD can be combined to recover the unavailable or missing data blocks using a simple hardware exclusive-OR procedure. Such systems often employ standby ("hot") DASDs that may be switched into the DASD array for immediate reconstitution of the data frozen in a failed DASD.

Current US Original Classification - CCOR (1):

711/114

US-PAT-NO: 6098146

DOCUMENT-IDENTIFIER: US 6098146 A
See image for Certificate of Correction

TITLE: Intelligent backplane for collecting and reporting information in an SSA system

-------- KWIC --------

Detailed Description Text - DETX (27):
As before, the bus transceiver 69 enables the microcontroller 70 to transmit commands to control register 68-1 and drive write protect register 68-2 via the bus 49. The control register 68-1 is a bay power control register. Each disk bay includes a power control line coupled to a corresponding control bit of the bay power control register 68-1. When the control bit is set to a logical "0", the corresponding power control line is de-asserted and the drive bay is powered down. If, however, the control bit is set to a logical "1", the power control line is asserted and the drive bay is powered up. Thus, by setting or resetting selected control bits of the bay power control register 68-1, the SSA initiator may selectively power up or power down a selected disk bay. For example, when a drive is inserted or "hot plugged" into a selected bay, the microcontroller 70 will detect the change in state in the corresponding status bit of the drive presence register 66-1 and notify the SSA initiator that a drive is being inserted in a selected disk bay. In response thereto, the SSA initiator powers up the bays into which the drive is inserted. Since the drive bay into which the drive is inserted is powered down at the time of insertion, hot plugging a drive into this bay does not affect the power provided to the other drives inserted in the disk pod.

Current US Cross Reference Classification - CCXR (1):
711/114

DOCUMENT-IDENTIFIER: US 20030014587 A1

TITLE: Intelligent backplane for serial storage architectures method and system

-------- KWIC -------

Current US Classification, US Primary Class/Subclass - CCPR (1):
711/114

Detail Description Paragraph - DETX (26):
[0042] As before, the bus transceiver 69 enables the microcontroller 70 to transmit commands to control register 68-1 and drive write protect register 68-2 via the bus 49. The control register 68-1 is a bay power control register. Each disk bay includes a power control line coupled to a corresponding control bit of the bay power control register 68-1. When the control bit is set to a logical "0", the corresponding power control line is de-asserted and the drive bay is powered down. If, however, the control bit is set to a logical "1", the power control line is asserted and the drive bay is powered up. Thus, by setting or resetting selected control bits of the bay power control register 68-1, the SSA initiator may selectively power up or power down a selected disk bay. For example, when a drive is inserted or "hot plugged" into a selected bay, the microcontroller 70 will detect the change in state in the corresponding status bit of the drive presence register 66-1 and notify the SSA initiator that a drive is being inserted in a selected disk bay. In response thereto, the SSA initiator powers up the bays into which the drive is inserted. Since the drive bay into which the drive is inserted is powered down at the time of insertion, hot plugging a drive into this bay does not affect the power provided to the other drives inserted in the disk pod.

DOCUMENT-IDENTIFIER: US 20030014598 A1

TITLE: Software raid methods and apparatuses including server usage based write delegation

--------- KWIC --------

Current US Classification, US Secondary Class/Subclass - CCSR (1):
711/114

Detail Description Paragraph - DETX (20):
[0032] Continuing to refer to FIG. 1, at any instant in time, each server 10.sub.i is running at some percent of capacity, also known as its usage level or just usage. This is depicted by the vertical bar graph 100.sub.i. Server 10.sub.1 is shown running at 90%, server 10.sub.i is shown running at 10%, and server 10.sub.N is shown running at 50%. These numbers can vary chaotically with time, and it is typical for some servers to be more heavily used on an average basis than others. It's also typical to have standby nodes in the cluster that are idle most of the time.

US-PAT-NO: 6078990

DOCUMENT-IDENTIFIER: US 6078990 A

TITLE: Volume set configuration using a single operational view

-------- KWIC --------

Detailed Description Text - DETX (107):

SCSILib 514 provides an interface to allow SCSI command descriptor blocks (CDBs) other than normal reads and writes to be sent to devices. Through this interface, requests like Start and Stop Unit will be used to spin and spin down disks, and Send and Receive Diagnostics will be used to monitor and control enclosure devices. All SCSILib routines are synchronous. The returning of the called function indicates the completion of the request.

Current US Original Classification - CCOR (1):

711/114

US-PAT-NO: 5666538

DOCUMENT-IDENTIFIER: US 5666538 A

TITLE: Disk power manager for network servers

-------- KWIC --------

Brief Summary Text - BSTX (17):

A further aspect of the present invention is a method of managing disk drive power within a network server having a processor which executes commands contained in a plurality of application modules and having a plurality of disk drives (disk drive subsystem) responsive to the commands executed by the processor. The method comprises the steps of loading a device driver application module to be executed by the processor; loading a monitoring application module to be executed by the processor; loading a console application module to be executed by the processor; monitoring accesses to the disk drives by the monitoring application module to generate a statistical record of the accesses to the disk drives; inputting data relating to spin-up and spin-down of the disk drives via the console application module based upon the statistical record of the accesses to the disk drives; and controlling spin-ups and spin-downs of the disk drives to reduce power consumption of the disk drives via commands executed by the processor to the device driver module based upon the input data.

Claims Text - CLTX (18):

controlling spin-ups and spin-downs of said disk drives to reduce power consumption of said disk drives via commands executed by said processor to said device driver module based upon said input data, said commands based upon the time of day and day of the week.

US-PAT-NO: 5961613

DOCUMENT-IDENTIFIER: US 5961613 A

TITLE: Disk power manager for network servers

-------- KWIC --------

Brief Summary Text - BSTX (17):

A further aspect of the present invention is a method of managing disk drive power within a network server having a processor which executes commands contained in a plurality of application modules and having a plurality of disk drives (disk drive subsystem) responsive to the commands executed by the processor. The method comprises the steps of loading a device driver application module to be executed by the processor; loading a monitoring application module to be executed by the processor; loading a console application module to be executed by the processor; monitoring accesses to the disk drives by the monitoring application module to generate a statistical record of the accesses to the disk drives; inputting data relating to spin-up and spin-down of the disk drives via the console application module based upon the statistical record of the accesses to the disk drives; and controlling spin-ups and spin-downs of the disk drives to reduce power consumption of the disk drives via commands executed by the processor to the device driver module based upon the input data.

server 110). As
s first loaded
nitor module 1120
10. Finally,
nted within an

(42) In accordance with a further aspect of the present invention, shown in FIG. 9, failure analysis can be performed using the histogram feature of the invention. In particular, once the histogram is generated, as indicated in an activity block 1300, the network administrator may observe long-term trends in accesses to individual disk drives 240, within the disk subsystem 245, as indicated within an activity block 1310. These long-term trends may indicate that certain disk drives are being accessed significantly more often than other disk drives. In order to reduce wear on the disk drives which are accessed more frequently, the network administrator may redistribute the information stored on the disk drives within the disk subsystem 245 so that information which is accessed more often is moved to disk drives with a history of fewer disk accesses. The redistribution of information stored on individual disk drives is represented in an activity block 1320 of FIG. 9. This causes accesses to the disk drives to be distributed more evenly so that any particular disk drive is less likely to fail over long-term use (i.e., because a disk drive is more likely to fail if the disk drive is accessed more often).

(43) It should be noted here that, although the present invention can be implemented in network servers which are configured to run various network operating systems (e.g., NOVELL NETWARE), certain network operating systems require that accesses to the disk drives are made every few minutes, even while no users are logged onto the network. Thus, such systems may not derive the full advantages afforded by the disk power manager of the present invention. In one advantageous configuration, depicted in FIG. 10, the network server 110 is advantageously configured to include a small disk drive 242 separate from the other disks in the disk subsystem 245. The small disk drive 242 primarily contains data which are regularly accessed by the operating system, while the other disks in the disk subsystem 245 store data which are not accessed by the operating system during down-time (i.e., when no users are logged onto the network). Thus, only the small disk drive 242 is continuously spun-up, while the other disk drives are allowed to spin down at the designated intervals. Thus, the disk drives which are not accessed by the operating system can be spun-down to provide additional energy savings.

(44) Although the preferred embodiment has been described and illustrated above, those skilled in the art will appreciate that various changes and obvious modifications do not depart from the spirit or essence of the invention. For example, alarm set points based upon the number of disk accesses (i.e., where an alarm sounds to alert the network administrator once a certain number of accesses has been exceeded) could be included within a preferred embodiment of the invention. Furthermore, alarm indications could be made using a "pop-up" error message screen, by sending an E-mail, or by activating a pager. In addition, the disk power manager could easily be adapted for use with redundant arrays of inexpensive disks (RAIDs). Accordingly, the scope of the present invention is limited only by the language

RAID

US-PAT-NO: 6192481

DOCUMENT-IDENTIFIER: US 6192481 B1

TITLE: Structure and method for power sequencing of disk drives in a computer system

-------- KWIC --------

Abstract Text - ABTX (1):

A method and device for detecting and handling non-responsive devices in a computer system where the device non-responsiveness may be due to a powered-down status rather than a device failure, and more particularly to such computer systems when the devices are RAID disk drives. By scanning all devices connected in a configuration and maintaining a count of devices that time out without responding, a determination can be made as to whether the devices are powered off or are experiencing some other problem that requires attention of a system support technician.

Brief Summary Text - BSTX (2):

This invention relates generally to apparatus and method for handling a non-responsive device in a computer system where the device non-responsiveness may be due to a powered-down status and not a device failure, and more particularly to such computer systems when the devices are RAID disk drives.

Brief Summary Text - BSTX (5):

This situation is particularly an issue in host computer systems which serve as database or information servers, and which typically have a host computer and one or more racks or shelves of rotating disk drive storage devices for storing the information. Customarily, each of the host computer processor rack and disc drive racks are powered by separate switchable power supplies. Unfortunately, the order and timing of the power up and power-down of the several racks effects the start-up or boot routine at system initialization, and may cause an error condition on shut-down or power-off.

Brief Summary Text - BSTX (6):

These conditions have been tolerated in the past by (I) indoctrinating personnel as to the proper power-up and power-down sequence for the host computer and attached devices, (ii) by providing a master power-on switch for all of the equipment, or (iii) by correcting corrupted or erroneous device or system configuration files after the problem has occurred. Unfortunately, neither of the first two options has been entirely successful so that corruption still occurs, and when such corruption occurs, correction typically requires the intervention of a skilled computer administrator.

US-PAT-NO: 6449725

DOCUMENT-IDENTIFIER: US 6449725 B2

TITLE: Method and computer program product for diagnosing and handling non-responsive device in a computer system

--------- KWIC --------

Abstract Text - ABTX (1):

A method and device for detecting and handling non-responsive devices in a computer system where the device non-responsiveness may be due to a powered-down status rather than a device failure, and more particularly to such computer systems when the devices are RAID disk drives. By scanning all devices connected in a configuration and maintaining a count of devices that time out without responding, a determination can be made as to whether the devices are powered off or are experiencing some other problem that requires attention of a system support technician.

Brief Summary Text - BSTX (2):

This invention relates generally to apparatus and method for handling a non-responsive device in a computer system where the device non-responsiveness may be due to a powered-down status and not a device failure, and more particularly to such computer systems when the devices are RAID disk drives.

Brief Summary Text - BSTX (5):

This situation is particularly an issue in host computer systems which serve as database or information servers, and which typically have a host computer and one or more racks or shelves of rotating disk drive storage devices for storing the information. Customarily, each of the host computer processor rack and disc drive racks are powered by separate switchable power supplies. Unfortunately, the order and timing of the power up and power-down of the several racks effects the start-up or boot routine at system initialization, and may cause an error condition on shut-down or power-off.

Brief Summary Text - BSTX (6):

These conditions have been tolerated in the past by (I) indoctrinating personnel as to the proper power-up and power-down sequence for the host computer and attached devices, (ii) by providing a master power-on switch for all of the equipment, or (iii) by correcting corrupted or erroneous device or system configuration files after the problem has occurred. Unfortunately, neither of the first two options has been entirely successful so that corruption still occurs, and when such corruption occurs, correction typically requires the intervention of a skilled computer administrator.

US-PAT-NO: 5403639

DOCUMENT-IDENTIFIER: US 5403639 A

TITLE: File server having snapshot application data groups

-------- KWIC --------

Detailed Description Text - DETX (21):

Each of the disk drives 122-1 to 125-r in disk drive subset 103-1 can be considered a disk subsystem that consists of a disk drive mechanism and its surrounding control and interface circuitry. The disk drive consists of a commodity disk drive which is a commercially available hard disk drive of the type that typically is used in personal computers. A control processor associated with the disk drive has control responsibility for the entire disk drive and monitors all information routed over the various serial data channels that connect each disk drive 122-1 to 125-r to control and drive circuits 121. Any data transmitted to the disk drive over these channels is stored in a corresponding interface buffer which is connected via an associated serial data channel to a corresponding serial/parallel converter circuit. A disk controller is also provided in each disk drive to implement the low level electrical interface required by the commodity disk drive. The commodity disk drive has an interface which must be interfaced with control and drive circuits 121. The disk controller provides this function. Disk controller provides serialization and deserialization of data, CRC/ECC generation, checking and correction and NRZ data encoding. The addressing information such as the head select and other type of control signals are provided by control and drive circuits 121 to commodity disk drive 122-1. This communication path is also provided for diagnostic and control purposes. For example, control and drive circuits 121 can power a commodity disk drive down when the disk drive is in the standby mode. In this fashion, commodity disk drive remains in an idle state until it is selected by control and drive circuits 121.

Current US Cross Reference Classification - CCXR (3):

711/114

US-PAT-NO: 5193184

DOCUMENT-IDENTIFIER: US 5193184 A

TITLE: Deleted data file space release system for a dynamically mapped virtual data storage subsystem

-------- KWIC -------

Detailed Description Text - DETX (10):

Each of the disk drives 122-1 to 125-r in disk drive subset 103-1 can be considered a disk subsystem that consists of a disk drive mechanism and its surrounding control and interface circuitry. The disk drive consists of a commodity disk drive which is a commercially available hard disk drive of the type that typically is used in personal computers. A computer processor associated with the disk drive has control responsibility for the entire disk drive and monitors all information routed over the various serial data channels that connect each disk drive 122-1 to 125-r to control and drive circuits 121. Any data transmitted to the disk drive over these channels is stored in a corresponding interface buffer which is connected via an associated serial data channel to a corresponding serial/parallel converter circuit. A disk controller is also provided in each disk drive to implement the low level electrical interface required by the commodity disk drive. The commodity disk drive has an ESDI interface which must be interfaced with control and drive circuits 121. The disk controller provides this function. Disk controller provides serialization and deserialization of data, CRC/ECC generation, checking and correction and NRZ data encoding. The addressing information such as the head select and other type of control signals are provided by control and drive circuits 121 to commodity disk drive 122-1. This communication path is also provided for diagnostic and control purposes. For example, control and drive circuits 121 can power a commodity disk drive down when the disk drive is in the standby mode. In this fashion, commodity disk drive remains in an idle state until it is selected by control and drive circuits 121.

Current US Cross Reference Classification - CCXR (2):

711/114

US-PAT-NO: 5155835

DOCUMENT-IDENTIFIER: US 5155835 A

TITLE: Multilevel, hierarchical, dynamically mapped data storage subsystem

-------- KWIC --------

Detailed Description Text - DETX (16):

Each of the disk drives 122-1 to 125-r in disk drive subset 103-1 can be considered a disk subsystem that consists of a disk drive mechanism and its surrounding control and interface circuitry. The disk drive consists of a commodity disk drive which is a commercially available hard disk drive of the type that typically is used in personal computers. A control processor associated with the disk drive has control responsibility for the entire disk drive and monitors all information routed over the various serial data channels that connect each disk drive 122-1 to 125-r to control and drive circuits 121. Any data transmitted to the disk drive over these channels is stored in a corresponding interface buffer which is connected via an associated serial data channel to a corresponding serial/parallel converter circuit. A disk controller is also provided in each disk drive to implement the low level electrical interface required by the commodity disk drive. The commodity disk drive has an ESDI interface which must be interfaced with control and drive circuits 121. The disk controller provides this function. Disk controller provides serialization and deserialization of data, CRC/ECC generation, checking and correction and NRZ data encoding. The addressing information such as the head select and other type of control signals are provided by control and drive circuits 121 to commodity disk drive 122-1. This communication path is also provided for diagnostic and control purposes. For example, control and drive circuits 121 can power a commodity disk drive down when the disk drive is in the standby mode. In this fashion, commodity disk drive remains in an idle state until it is selected by control and drive circuits 121.

Current US Original Classification - CCOR (1):

711/114

US-PAT-NO: 5566316

DOCUMENT-IDENTIFIER: US 5566316 A

TITLE: Method and apparatus for hierarchical management of data storage elements in an array storage device

-------- KWIC --------

Detailed Description Text - DETX (6):

Each of the data storage elements D11-LS12 in data storage module 211 can be implemented by a disk subsystem that consists of a disk drive mechanism and its surrounding control and interface circuitry. The disk drive consists of a commodity disk drive which can be a commercially available hard disk drive of the type that typically is used in personal computers. A control processor associated with the disk drive has control responsibility for the entire disk drive and monitors all information routed over the various serial data channels that connect each disk drive to drive circuits. Any data transmitted to the disk drive over these channels is stored in a corresponding interface buffer which is connected via an associated serial data channel to a corresponding serial/parallel converter circuit. A disk controller is also provided in each disk drive to implement the low level electrical interface required by the commodity disk drive. The commodity disk drive has a standard interface which must be interfaced with drive circuits. The disk controller provides this function. Disk drive controller provides serialization and deserialization of data, CRC/ECC generation, checking and correction and data encoding. The addressing information such as the head select and other type of control signals are provided by drive circuits to commodity disk drive D11. This communication path is also provided for diagnostic and control purposes. For example, drive circuits can power a commodity disk drive down when the disk drive is in the standby mode. In this fashion, commodity disk drive remains in an idle state until it is selected by drive circuits.

Current US Original Classification - CCOR (1):

711/114

US-PAT-NO: 5581724

DOCUMENT-IDENTIFIER: US 5581724 A

TITLE: Dynamically mapped data storage subsystem having multiple open destage cylinders and method of managing that subsystem

--------- KWIC --------

Detailed Description Text - DETX (15):

Each of the disk drives 122-1 to 125-r in disk drive array 103-1 can be considered a disk subsystem that consists of a disk drive mechanism and its surrounding control and interface circuitry. The disk drive consists of a commodity disk drive which is a commercially available hard disk drive of the type that typically is used in personal computers. A control processor associated with the disk drive has control responsibility for the entire disk drive and monitors all information routed over the various serial data channels that connect each disk drive 122-1 to 125-r to control and drive circuits 121. Any data transmitted to the disk drive over these channels is stored in a corresponding interface buffer which is connected via an associated serial data channel to a corresponding serial/parallel converter circuit. A disk controller is also provided in each disk drive to implement the low level electrical interface required by the commodity disk drive. The commodity disk drive has an ESDI interface which must be interfaced with control and drive circuits 121. The disk controller provides this function. Disk controller provides serialization and deserialization of data, CRC/ECC generation, checking and correction and NRZ data encoding. The addressing information such as the head select and other type of control signals are provided by control and drive circuits 121 to commodity disk drive 122-1. This communication path is also provided for diagnostic and control purposes. For example, control and drive circuits 121 can power a commodity disk drive down when the disk drive is in the standby mode. In this fashion, commodity disk drive remains in an idle state until it is selected by control and drive circuits 121.

Current US Original Classification - CCOR (1):

711/114

US-PAT-NO: 6529995

DOCUMENT-IDENTIFIER: US 6529995 B1

TITLE: Method and apparatus for maintaining and restoring mapping table entries and data in a raid system

-------- KWIC --------

Detailed Description Text - DETX (19):

Each of the disk drives 122-1 to 125-r in disk drive subset 103-1 can be considered a disk subsystem that consists of a disk drive mechanism and its surrounding control and interface circuitry. The disk drive consists of a commodity disk drive which is a commercially available hard disk drive of the type that typically is used in personal computers. A control processor associated with the disk drive has control responsibility for the entire disk drive and monitors all information routed over the various serial data channels that connect each disk drive 122-1 to 125-r to control and drive circuits 121. Any data transmitted to the disk drive over these channels is stored in a corresponding interface buffer which is connected via an associated serial data channel to a corresponding serial/parallel converter circuit. A disk controller is also provided in each disk drive to implement the low level electrical interface required by the commodity disk drive. The commodity disk drive has an ESDI interface which must be interfaced with control and drive circuits 121. The disk controller provides this function. Disk controller provides serialization and deserialization of data, CRC/ECC generation, checking and correction and NRZ data encoding. The addressing information such as the head select and other type of control signals are provided by control and drive circuits 121 to commodity disk drive 122-1. This communication path is also provided for diagnostic and control purposes. For example, control and drive circuits 121 can power a commodity disk drive down when the disk drive is in the standby mode. In this fashion, commodity disk drive remains in an idle state until it is selected by control and drive circuits 121.

Current US Original Classification - CCOR (1):

711/114

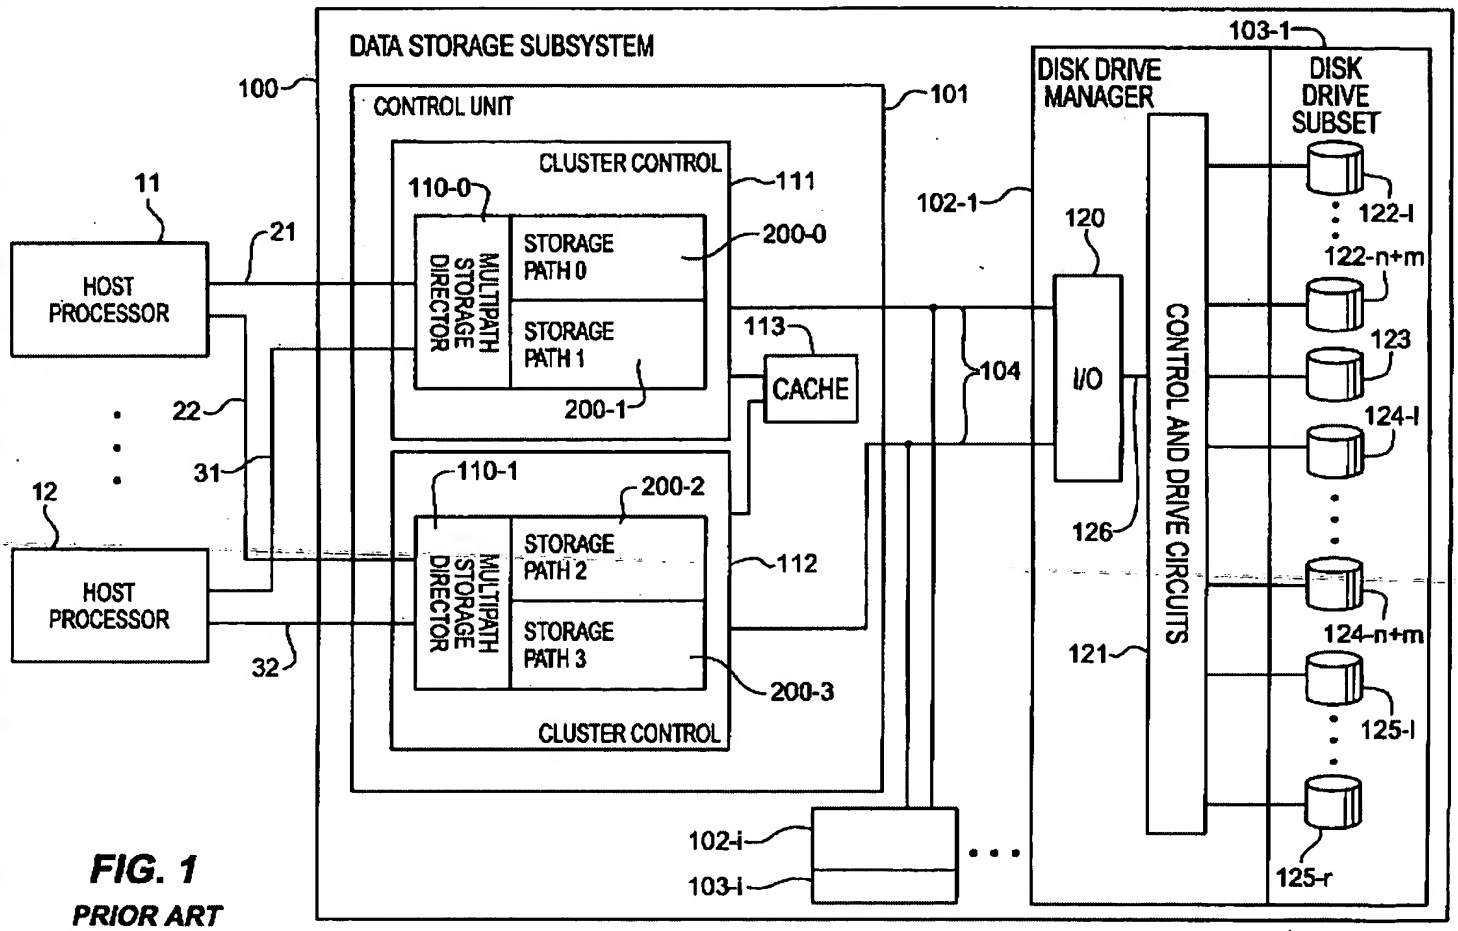

FIG. 1
PRIOR ART

DOCUMENT-IDENTIFIER: US 20030200473 Á1

TITLE: System and method for activity or event based dynamic energy conserving server reconfiguration

-------- KWIC --------

Detail Description Paragraph - DETX (245):

[0280] In the RAID 1 (and RAID 10) configurations, only one drive (primary or mirror) or one set of drives (primary set or mirror set) need to be available or powered ON at a time to support such read operations. The identical nature of the data stored on the primary and mirror drives only changes in response to a write operation and is unchanged by the more frequent read operations. One of the drives can be either in a standby mode (such as a mode where power is provided by the spindle motor is not operating) or with operating power (e.g. operating voltage and/or current) completely removed. Various shades of power conservation may be applied between completely on and completely off, particularly if the internal circuitry of the disc drive and control electronics and any on-board buffer memory or the like are designed with staged power conservation features. It is noted that since the primary and secondary disc drives store identical data and are completely interchangeable from a functional (and likely from a physical standpoint) there is little need to identify that it is the primary or the secondary drive that is powered off or placed into a standby mode, reduced power consumption mode, power conservation mode, or simply powered off. More accurately, we may refer to the drives as the active drive (or active drive set) and the inactive drive (or inactive drive set).

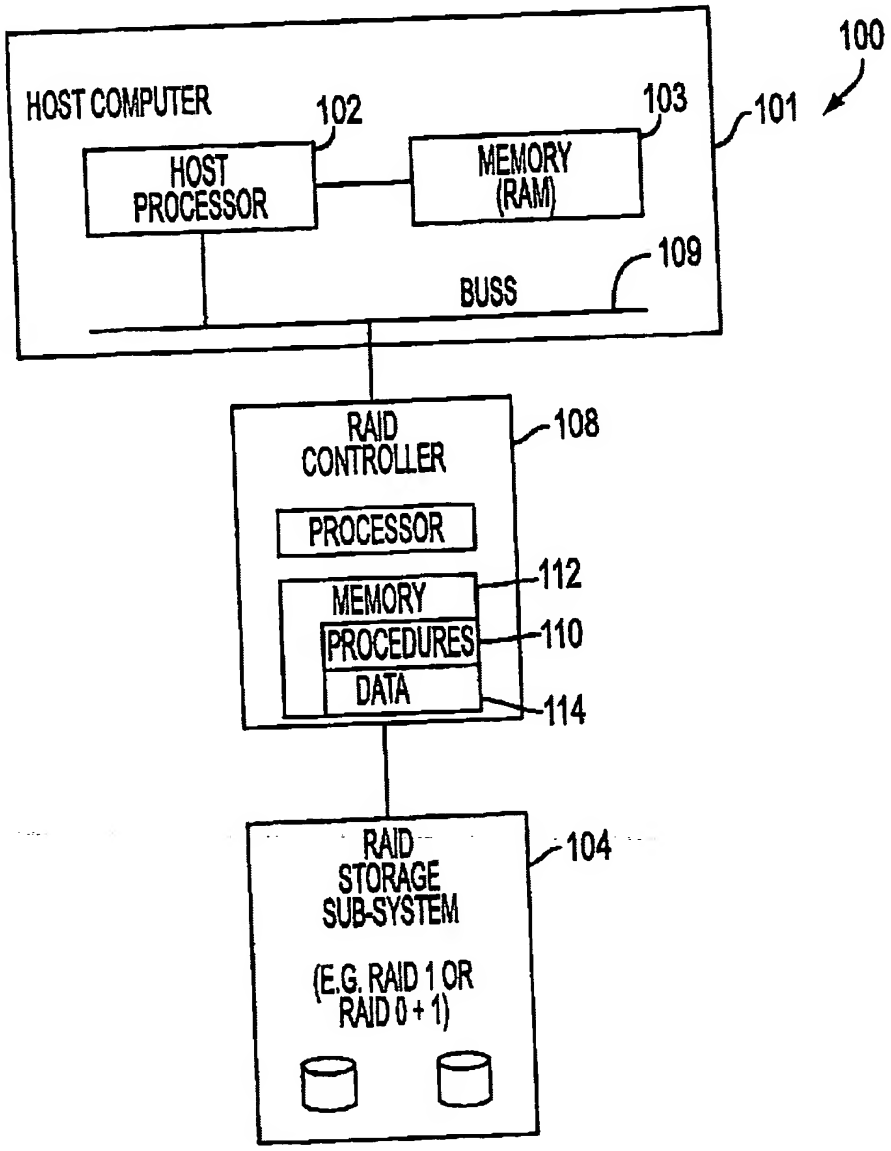

FIG. 25

DOCUMENT-IDENTIFIER: US 20030188208 A1

TITLE: System, method, and architecture for dynamic server power management and dynamic workload management for multi-server environment

-------- KWIC --------

Detail Description Paragraph - DETX (245):

[0280] In the RAID 1 (and RAID 10) configurations, only one drive (primary or mirror) or one set of drives (primary set or mirror set) need to be available or powered ON at a time to support such read operations. The identical nature of the data stored on the primary and mirror drives only changes in response to a write operation and is unchanged by the more frequent read operations. One of the drives can be either in a standby mode (such as a mode where power is provided by the spindle motor is not operating) or with operating power (e.g. operating voltage and/or current) completely removed. Various shades of power conservation may be applied between completely on and completely off, particularly if the internal circuitry of the disc drive and control electronics and any on-board buffer memory or the like are designed with staged power conservation features. It is noted that since the primary and secondary disc drives store identical data and are completely interchangeable from a functional (and likely from a physical standpoint) there is little need to identify that it is the primary or the secondary drive that is powered off or placed into a standby mode, reduced power consumption mode, power conservation mode, or simply powered off. More accurately, we may refer to the drives as the active drive (or active drive set) and the inactive drive (or inactive drive set).

EAST Browser - (11) (((disk or d... | US 5961613 A | Tag: S,T1 | Doc: 9/11

File Edit View Tools Window Help

	Document ID	Kind Codes	Source	Issue Date	Pages	Title
1	US 20030200473 A1		US-PGPUB	20031023	77	System and method for activity or event bas
2	US 20030188208 A1		US-PGPUB	20031002	79	System, method, and architecture for dyna
3	US 20020144057 A1		US-PGPUB	20021003	14	Archival data storage system and method
4	US 20020062454 A1		US-PGPUB	20020523	73	Dynamic power and workload management
5	US 20020007464 A1		US-PGPUB	20020117	75	Apparatus and method for modular dynamic
6	US 20020004915 A1		US-PGPUB	20020110	75	System, method, architecture, and compute
7	US 20020004913 A1		US-PGPUB	20020110	74	Apparatus, architecture, and method for int
8	US 20020004912 A1		US-PGPUB	20020110	74	System, architecture, and method for logica
9	US 5961613 A		USPAT	19991005	24	Disk power manager for network servers
10	US 5666538 A		USPAT	19970909	25	Disk power manager for network servers
11	US 5448719 A		USPAT	19950905	30	Method and apparatus for maintaining and r

all same spec

Details Text Image HTML

US 20030200473A1

(19) **United States**

(12) **Patent Application Publication** (10) **Pub. No.: US 2003/0200473 A1**

Fung (43) **Pub. Date: Oct. 23, 2003**

(54) **SYSTEM AND METHOD FOR ACTIVITY OR EVENT BASED DYNAMIC ENERGY CONSERVING SERVER RECONFIGURATION**

(75) Inventor: **Henry T. Fung**, San Jose, CA (US)

Correspondence Address:
R. Michael Ananian
FLEHR HOHBACH TEST
ALBRITTON & HERBERT LLP
Four Embarcadero Center, Suite 3400
San Francisco, CA 94111 (US)

(73) Assignee: **Amphus, Inc.**

(21) Appl. No.: **09/860,995**

(22) Filed: **May 18, 2001**

Related U.S. Application Data

(63) Continuation of application No. 09/121,352, filed on Jul. 23, 1998, now Pat. No. 6,079,025.
Continuation of application No. 08/767,821, filed on Dec. 17, 1996, now Pat. No. 5,892,959.
Continuation of application No. 08/460,191, filed on Jun. 2, 1995, now abandoned.
Continuation of application No. 08/285,169, filed on Aug. 3, 1994, now abandoned.
Continuation of application No. 08/017,975, filed on Feb. 12, 1993, now Pat. No. 5,396,635.
Continuation of application No. 07/908,533, filed on Jun. 29, 1992, now abandoned.
Continuation of application No. 07/532,314, filed on Jun. 1, 1990, now abandoned.
Continuation of application No. 09/558,473, filed on Apr. 25, 2000, now Pat. No. 6,584,571.

(60) Provisional application No. 60/283,375, filed on Apr. 11, 2001. Provisional application No. 60/236,043, filed on Sep. 27, 2000. Provisional application No. 60/236,062, filed on Sep. 27, 2000.

Publication Classification

(51) **Int. Cl.**[7] **G06F 1/26; G06F 1/32**
(52) **U.S. Cl.** **713/320**

(57) **ABSTRACT**

Network architecture, computer system and/or server, circuit, device, apparatus, method, and computer program and control mechanism for managing power consumption and workload in computer system and data and information servers. Further provides power and energy consumption and workload management and control systems and architectures for high-density and modular multi-server computer systems that maintain performance while conserving energy and method for power management and workload management. Dynamic server power management and optional dynamic workload management for multi-server environments is provided by aspects of the invention. Modular network devices and integrated server system, including modular servers, management units, switches and switching fabrics, modular power supplies and modular fans and a special backplane architecture are provided as well as dynamically reconfigurable multi-purpose modules and servers. Backplane architecture, structure, and method that has no active components and separate power supply lines and protection to provide high reliability in server environment.